Infrastructure

Recall from part 1:

- Finder Chart gives you images from POSS, 2MASS, WISE, Spitzer (SEIP), AKARI, and IRAS in native pixel resolution (e.g., the pixels are the original ones from the data product from that survey). But it is limited by the size of the tiles provided by the corresponding survey.
- IRSA Viewer gives you images from many, many surveys at IRSA (and some not at IRSA), but it too is limited by the size of the tiles provided by the corresponding survey.
- For the surveys where you run off the tile rather dramatically (at least 2MASS), you can use Skyview to get a larger image. BUT you have to specify some parameters that you don't have to specify in Finder Chart or IRSA Viewer.
- Both ds9 and the IRSA tools allow you to overlay things. IRSA tools prefer IPAC table files (*.tbl files), but will accept regions files (*.reg); ds9 prefers regions files and doesn't know what to do with tbl files. Both regions and table files have unusual filename extensions, which makes Bob's email gag, and may make your computer complain too. They are both plain text files.
- The skills you're going to need here include: finding data, loading and viewing FITS images, aligning the images so that they all cover the same region of the sky, overlaying regions, changing the color stretch and maybe the color table. Also: moving back and forth between tools to find the capability you need.
- The reason I want you to pull the data from IRSA (or Skyview) directly and not just grab copies off the Box drive is so that you know how to find data after your work on this team is done.

Recall from our proposal:

- Our region is a cone of 0.5° radius, centered on 18:03:37.0 -24:23:12.

MY GOALS for you in doing this are to (1) develop a sense of what spatial resolution means and how it changes between telescopes, e.g., WISE vs. Spitzer vs. 2MASS vs. Herschel resolution; and (2) understand what the challenges will be for us in matching across wavelengths. (3) I also think we should, anticipating comments from the referees, work towards adding a column to our big table with the resolution of each dataset.

Ancillary goals (e.g., you can't do this worksheet without also accomplishing these): (1) get used to working with FITS files, manipulating stretches, etc.; (2) identifying objects in and measuring distances on FITS files; (3) learn how to use Finder Chart (and other IRSA tools), ds9, and Skyview as resources to be used down the road on whatever you find yourself doing next.

Quantitative measures of resolution.

Normally, to 'believe' a detection of anything, astronomers require that it be seen in more than 1 pixel. If something is seen in just 1 pixel, it's hard to tell if it's a single hot pixel, or a cosmic ray, or a real detection. Thus, spatial resolution, if cited without a "per pixel," is most frequently quoted as certainly more than 1 pixel, often ~2 pixels. What this physically means, in essence, is BOTH the following two questions: (1) "How many pixels have to be affected before I believe it is a real detection?" and (2) "How close do two sources have to be before I can no longer distinguish them as two individual sources?"

Real life numbers: the quoted resolution of IRAC is ~2 arcsec, but the native pixel size is 1.2 arcsec, and mosaics often have the pixels resampled to be 0.6 arcsec. The quoted resolution of the DSS is 1.7 arcsec per pixel (or about 2 arcsec, depending on the photographic plate). Most of the surveys and published catalogs we are using have been electronic from the start; however, the original POSS was (ground-based) photographs, so the spatial resolution was set by the seeing at Palomar that night, AND the size of the silver grains on the photographic emulsion. When the plates got scanned, during the digitization process, this got mapped into the pixels you see in the images.

Suggestions for comparing resolution qualitatively between surveys.

Create a three-color image using bands of your choice. Finder Chart used to let you do this, but that's not possible in the current version. IRSA Viewer is the best way do this using IRSA tools. On the main screen where you search for objects, you have a 3-color option, including loading from disk (as opposed to, say, IRSA holdings). However, IIRC, the final resolution of the 3-color image is set by whatever you load into the red plane, so even though your lowest spatial resolution observations are probably also the reddest, in this context, don't load them into the red plane. ds9 also allows you to make 3-color images. For ds9, you need to tell it, "Ok, I want to make a 3-color image now" (Frame/rgb) and then you can load in each plane separately (in the pop-up, pick the color plane, then do File/open. Change the color plane and go back to file/open, etc.). Regardless of tool, whichever image has the lowest spatial resolution should be super obvious, because, say, the sources will appear to have blue blobby rings if you load the lowest resolution image into the blue plane.

Align or align & blink images. You can flip through images in any of these tools; ds9 will do it for you if you want. For ds9, do file/open and find the first image; do frame/new then file/open and load the second image, etc. If you used the command line trick (ds9 *.fits), you will load all the images into individual tiles, in alphabetical order (which is most likely not wavelength order!). If you did them one-by-one, you will have them virtually in a stack, in the order you loaded them. To see all of them at once, click on 'frame' then 'tile.' To get it back to one at a time (in a virtual stack), pick 'single.' To line them up on the sky, pick from the top "frame" menu/match/frame/wcs to match them in terms of area on the sky. (That command means, "align all the images I have loaded in ds9 to be North up, all on the same spatial scale as the image I have selected when I initiate this command." WCS stands for world coordinate system, meaning that there is information about the ra, dec, and mapping of pixels to ra and dec in the FITS header.) To scroll through the whole stack, pick 'next' or 'previous', or go

ahead and blink them. You can configure the length of time spent on each frame. You can change the ordering - explore the menu options on the top "Frame" menu. In the 'single' frame case, the image you are looking at is the active one; in the 'tile' view, the one with the blue outline is the active one. Click on the tile to make it the active one. The commands are similar in the IRSA tools; there is a "match WCS" tickbox and you can view images one at a time or tiled in the tool.

Measuring resolutions

Since the resolution issues are going to be more important to us in the infrared, let's focus on those bands for most of our by-hand (as opposed to by-Google) work. Our infrared (through submm) inventory is:

- 2MASS
- Spitzer (SEIP, GLIMPSE)
- WISE
- SCUBA (partial coverage)
- Herschel (partial coverage)
- AKARI (only have images for FIS, not IRC)
- MSX
- IRAS

You should aim for, at minimum, empirically finding the resolution of 2MASS, Spitzer, and WISE. Use Finder Chart (or IRSA Viewer) to retrieve images of our region from at least 2MASS, Spitzer, and WISE. (If you want to do more, you will need also to use IRSA Viewer, rather than just Finder Chart.)

Q1.1 : Retrieve images of our area. For the images that it returns, what is the size of each pixel for each survey? (Option #1 to do this: Make the image big enough in your view of it that you can see pixels, and measure the size of it using ruler tools (not a real ruler). Option #2 to do this: look in the FITS header and find a useful keyword.) Try at least one image from each of the surveys.

Q1.2 : You will need to Google for this one. What is the original native pixel size for these surveys? Finder Chart gives you images that come straight from the original surveys, so they should match the original native pixel size for each survey.

Q1.3 : Are there any images you've retrieved that have "run off the edge" of a stored tile? (Hint: yes.) Which ones?

For the surveys where you have run off the tile rather dramatically (at least 2MASS), you can use Skyview to get a larger image. The four most important parameter choices Skyview gives you are:

- center position

- survey (wavelength)
- image size in pixels
- image size in degrees

Skyview will happily and without complaint or warning resample and regrid the pixels to whatever scale you want. What do you need to do to get 'native pixel' resolution out of Skyview? You should have the information from earlier questions to figure out how many pixels you need to cover our region, so go and do the math, and ask Skyview to give you a full-sized image of your desired size. Note that you can request more than one survey at a time, but Skyview will use the same parameters for each of them.

Q1.4 : Did you do the calculations right? Here's how to check. Look at the sources in the 2MASS image you retrieved from Finder Chart (which you know is native px size) and compare it to the sources in the image you retrieved from Skyview. Have you lost information? (To see what this looks like, try to make it lose information deliberately by asking for much larger pixels.) We'll work more with individual sources shortly.

Q1.5 : Skyview attempts to knit tiles together, but sometimes you can see the original tile boundaries, and it looks like a patchwork quilt. Do you see this here?

Q1.6 : For at least one frame from each of a few of the surveys we picked, from either your Finder Chart or Skyview images (assuming you are confident you have native pixel resolution), go and measure the sizes of 3 to 5 'typical' isolated point sources in these images. What kinds of sizes are you getting for each survey? (It is going to be hard to find 'typical' in IRAS; do what you can.) Changing the color table/stretch is useful for telling if the image is slightly asymmetric (implying a barely resolved companion) or saturated or other things.

Skyview won't give you Spitzer images, because Spitzer isn't an all-sky survey. But there are lots of large images available at IRSA from Spitzer. SEIP = Spitzer Enhanced Imaging Products, but this too works in tiles, and the request you give Finder Chart or IRSA Viewer may run off the edges of some of those tiles. There are data there, just not in the tile that the IRSA tools may be pulling for you. To find individual sources in regions off the tile it gives you, ask for a smaller region.

Inspecting Point Sources

Now we are going to start to look at the individual sources in our region. We're going to do this in spades soon… but let's start small.

Q2.1 : The IAU-compliant names of sources are based on positions. Many of the catalogs and papers that we have list some sort of unique ID within the survey, but its 'real' name is the position-based name, which is typically included in the catalogs if not all the journal articles (the journal articles are supposed to use position-based names, but they don't always). People often assign and use internal source IDs in papers because it's easier to say "source 346" in conversations with collaborators rather than the full phone number that might look

like 18033652-2423108. We will do something similar. But, why is it that IRAS sources are given as, e.g., "IRAS 18006-2422" and 2MASS sources are given as, e.g., "2MASS 18033652-2423108"?

Q2.2 : Here are five of the sources off our list of interesting objects. Do what you need to do to see if the source appears, and appears single, in all the IR bands to which we have access.
Hints: You are trying to decide if the same source appears at the nominal location of the source in all the bands ... and is apparently circular (like other point sources in that band, and comparably sized). You will probably have to use Finder Chart sometimes and IRSA Viewer sometimes. Find the images you need. Call them up in whatever FITS viewer works for you. Decide what size images you need, and/or how much you need to zoom. You may need to change the color stretch and/or color table. Are these sources present in all available bands in J band and longer? If not, which bands have this source? Are any bands saturated? The sources start easier and get harder.
Bonus: find the name of any counterparts.
Double-bonus: how far offset from the given position are the counterparts? (By the way, you don't need to calculate this by hand…)
Here is the list of sources. You can find these by hand, or, to make it easier, there is a regions file in the Box drive with these 5 positions (and labels).

1. 18:02:51.10 -24:16:56.8
2. 18:03:11.63 -24:11:56.8
3. 18:03:10.28 -24:26:59.5
4. 18:04:03.55 -24:50:45.6 – may be off the edge of the images you grabbed initially.
5. 18:03:01.83 -24:03:02.6 – note that this one may merit looking in MSX too.

The table in our proposal.

Now we need to work on the additional column in our proposal with the resolution of each of those surveys. You should be able to fill out several rows of this table now, based on the work above. For the other rows, especially the ones from journal articles, we should be able to get the resolution from reading the text in the paper, or the documentation that comes with the survey catalog. Go forth and fill out as many rows as you can!

Pulling it together.

Revisit the goals from above. Is this starting to make more sense?

Questions to be sure you know the answer to (for reference, I guess these are Q3.x):

1. How can you get access to data using Skyview? Using Finder Chart? When would you use one vs. the other?
2. Just because you have resampled an image to really tiny pixels, does it add information to the image? (Will you laugh at CSI and Law & Order and their compatriots when they wave a magic wand over an image? "Computer, ENHANCE!")

3. Will we see disks or rings in our data?
4. How does the spatial resolution compare among 2MASS, WISE, Spitzer/IRAC, Spitzer/MIPS, Herschel/PACS, etc.?
5. Is it possible that the sources seen as individual with MIPS-24 will break into pieces when viewed in the optical?
6. Is there any guarantee that a single source seen in any of our infrared data is really a single object?

Postscript on the resolution issues: Slight improvements are sometimes possible

By this point, I've hammered into you things about the native resolution from these various surveys. You should have a gut-level understanding now that you can't get more information out of the image than was recorded by it in the first place.

However.

I have swept some things under the rug. IRAS data was so interesting, and it was going to be so long before astronomers got any more data in those wavelengths on that scale, that very clever people got to work on how to get even more information out of IRAS data. Imagine those big IRAS pixels scanning over a patch of warm sky. The next time the spacecraft scans that same patch of sky, the pixels are offset a little bit from where it was on the last pass, and consequently the fluxes it measures are just a little different. Same for the next scan, and the next. If you have lots of scans over the same region, each of which are at slightly different positions (that's important), you can recover a little bit of the information on a slightly higher (better) spatial resolution. This page has some general information on the specific application of this method to IRAS, called "Hi-Res", along with example pictures. It uses the Maximum Correlation Method (MCM; H.H. Aumann, J.W. Fowler and M. Melnyk, 1990, AJ, 99, 1674). It is computationally expensive (meaning it takes a while to run), and requires lots of individual tweaking and customization, so it has not been run (blindly) over the whole sky. The degree of improvement is related to the number of scans; as for WISE, the number of passes is a function of the ecliptic latitude, so just running Hi-Res doesn't get you a specific improved resolution. Hi-Res got famous in the context of IRAS. People have developed ways to run this kind of algorithm on WISE and even Spitzer data, but we're not going to try and use it, as there are no particularly user-friendly interfaces to it (at least at the level we would need), and the incremental benefit we'd gain from this probably outweighs the work it would take to get there.

Note - critical to making this process work is that the camera moves between scans to slightly different positions, and the source it is looking at is not changing in brightness. Will this process work on security camera videos?

In the context of our project, we won't need to care about any of this, but I thought I should be complete in case anyone cares! :)